Curvature-Metric-Free Surface Remeshing via Principal Component Analysis

Yiqi Cai, Xiaohu Guo, Yang Liu, Wenping Wang, Weihua Mao, Zichun Zhong

Abstract—In this paper, we present a surface remeshing method with high approximation quality based on Principal Component Analysis. Given a triangular mesh and a user assigned polygon/vertex budget, traditional methods usually require the extra curvature metric field for the desired anisotropy to best approximate the surface, even though the estimated curvature metric is known to be imperfect and already self-contained in the surface. In our approach, this anisotropic control is achieved through the optimal geometry partition without this explicit metric field. The minimization of our proposed partition energy has the following properties: Firstly, on a $\mathbf{C}^2$ surface, it is theoretically guaranteed to have the optimal aspect ratio and cluster size as specified in approximation theory for L^∞ piecewise linear approximation. Secondly, it captures sharp features on practical models without any pre-tagging. We develop an effective merging-swapping framework to seek the optimal partition and construct polygonal/triangular mesh afterwards. The effectiveness and efficiency of our method are demonstrated through the comparison with other state-of-the-art remeshing methods.

Index Terms—Error-Driven, Principal Component Analysis, Generalized Variance, Approximation Theory.

arXiv:1510.03935v1 [cs.GR] 14 Oct 2015

1 Introduction

With the ever increasing use of 3D scanners, faithful triangular meshes with dense sampling rate are easy to get for a wide variety of applications. However, for practical purposes, using a huge number of triangles as the surface representation becomes a burden for displaying, storing and processing. Thus a mesh-coarsening step introducing minimal geometric error becomes a quest for most real applications. That is, given a triangular mesh and a finite user assigned polygon/vertex budget, what is the best piecewise-linear triangulation approximation of the surface with respect to the L^2 or L^∞ metric? Despite the fundamental nature of this problem, it has been pointed out to be an NP-hard problem [1]. Due to its complexity, the optimality pursued by recent approaches refers to the optimal asymptotic behavior of the infinitesimal triangles given in approximation theory [2] [3] [4] (see Sec. 2).

Previous approaches can be classified into anisotropic surface remeshing approaches and the error-driven approaches. To best approximate the surface, anisotropic surface remeshing methods rely on the curvature metric to serve as the Riemannian metric to specify the desired anisotropy. However, an accurate curvature estimation is known to be difficult on discrete meshes, especially at feature regions, since it requires a careful definition of local neighborhood. More specifically, to estimate the curvature at the mesh vertices, usually the Euclidean k-nearest neighbors [5] or the ring-neighborhood [6] are selected to perform local fitting. Therefore, recent anisotropic mesh generation work [5] [7] adopt a smoothing step as preprocessing to reduce the impact of the noises and outliers arising from the estimation process. Instead of explicitly extracting the curvature information from the surface, the error-driven approach tackles the surface approximation problem from a different perspective: finding a partition that clusters mesh primitives with similar geometric characteristics, which seems to be a natural way for mesh coarsening [8]. The original NP-hard problem is casted as the problem to seek optimal partition of the original surface according to the partition error [8] [9], which is usually the summation over the fitting error from each partition region's *shape proxy* [8] to its underlying primitives. We discuss the limitations of previous partition error definitions in Sec. 3.

We solve the surface approximation problem via the optimal partition problem and consider our method as one of the error-driven methods. Different from Quadric Error Metrics (QEM) [9] and Variational Shape Approximation (VSA) [8], a partition energy based on the Principal Component Analysis (PCA) is proposed. The minimization of our proposed partition energy has the following properties: Firstly, on a $\mathbf{C}^2$ surface, it is theoretically guaranteed to have the optimal aspect ratio and cluster size asymptotically as specified in approximation theory for L^∞ piecewise linear approximation, without any explicit estimation of curvature metric tensors. Secondly, it captures sharp features on practical models without any pre-tagging. An efficient merging-swapping framework is proposed to seek the optimal partition and allows the processing of meshes that have up to tens of millions of triangular faces easily performed on a desktop. The polygonal/triangular mesh is constructed afterwards. The main contributions of our work include the following aspects:

- We propose a partition energy based on shape variance analysis whose optimization automatically

- *Y. Cai and X. Guo are with the Department of Computer Science, The University of Texas at Dallas, Richardson, TX 75083.*
- *Y. Liu is with Microsoft Research Asia, Beijing, China, 100080.*
- *W. Wang is with the Department of Computer Science, The University of Hong Kong, Hong Kong, China.*
- *W. Mao is with the Department of Radiation Oncology, The University of Texas Southwestern Medical Center, Dallas, TX 75390.*
- *Z. Zhong is with the Department of Computer Science, Wayne State University, Detroit, MI 48202.*
- *Corresponding Author: X. Guo, Email: xguo@utdallas.edu*

leads to the optimal clustering, i.e. each cluster has the optimal aspect ratio and relative size to approximate the surface in the asymptotical sense.
- We present an efficient discrete merging-swapping algorithm to compute the optimal partition. To compute the change of the partition energy arising from the proposed merging and swapping operations, only several float number multiplications are needed.

2 Previous Anisotropic Surface Meshing Approaches

Anisotropic surface meshing has been widely used as a surface modeling technique to serve the remeshing purpose. Its optimality of piecewise-linear approximation has been proven in approximation theory [4] that with respect to the general L^p metric, the optimal mesh is obtained when each triangle is elongated along the direction of the eigenvectors of its curvature tensor and the error is equidistributed on each triangle.

Given the set of sites $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^n$ to sample the surface, Centroidal Voronoi Tessellation (CVT) [10] is a special Voronoi Diagram (VD) [11] that each site $\mathbf{x}_i$ coincides with the centroid of its Voronoi cell. When a Riemannian metric tensor field is endowed to measure the way distance and angle are distorted, CVT is generalized to Anisotropic CVT (ACVT) [12] to take the anisotropy into account. From the perspective of energy optimization, there are two widely-used approaches for CVT/ACVT computation: the Lloyd relaxation [13] and the quasi-Newton energy optimization scheme [14].

The main challenge to apply the ACVT concept on surface remeshing is the time-consuming computation of the Anisotropic Voronoi Diagram (AVD). One effective way to avoid the complicated computation is to transform the anisotropic surface meshing problem to an isotropic one in a higher dimensional space. Lévy and Bonneel [15] proposed to transform the anisotropic meshing problem on 3D surface to the isotropic one on 6D surface, which can be solved efficiently by CVT with Voronoi Parallel Linear Enumeration. However, their method does not provide users the flexibility to control the Riemannian metric tensor field. To overcome this problem, Zhong et al. [16] incorporated the metric tensor into the formulation of their inter-particle energy and optimized the energy in a higher dimensional isotropic "embedding space". Another approach to accelerate AVD computation is the discrete clustering approach proposed by Valette et al. [6]. This discrete ACVT is fast to compute since the ACVT energy of the discrete Voronoi diagrams can be evaluated efficiently.

Other methods model the anisotropic meshing problem from different perspectives. Chen et al. [4] [17] proposed Optimal Delaunay triangulation (ODT) to acquire the optimal anisotropic mesh that minimizes the difference between a function and its interpolation over the mesh. Fu et al. [7] extended ODT to Locally Convex Triangulation (LCT) for the cases where the given Riemannian metric is not the Hessian of a global function. Delaunay refinement methods [5] [18] [19] insert points in Delaunay triangution to refine local mesh quality. The final anisotropic meshing usually guarantees certain measures of mesh quality but does not provide a globally optimal result.

Anisotropic surface meshing methods usually generate high-quality meshes w.r.t. the input Riemannian metric. However, when it is used as a surface remeshing tool with the curvature tensor serves as the Riemannian metric, the inaccuracy arising from the curvature estimation step inevitably results in an impairment of the surface approximation. Like the error-driven approaches, our method does not require users to provide any explicit curvature metric. Nevertheless, our partition result forms like implicit AVDs, i.e. the anisotropy naturally agrees with the implicitly contained curvature information. The formal proof of the optimality is shown in Sec. 4. Our approach utilizes the discrete clustering method [6] during variational partition. But unlike Valette et al.'s approach [6], our method adopts a different energy formulation that does not require any explicit curvature estimation.

3 Limitations of Previous Error-Driven Approaches

The error-driven approaches cast the problem of surface approximation to the problem of geometric partition. When geometric primitives with similar shape characteristics are classified together, it forms a simplified representation of the original shape.

Previous methods generally utilize shape proxy [8] to guide the geometric partition process. Given the surface $\mathbf{S}$ and the user-assigned partition number n, the goal is to find the optimal partition (a partition is the set of clusters $\{\mathbf{C}_i\}_{i=1}^n$ that satisfy $\mathbf{C}_i \cap \mathbf{C}_j = \emptyset$ for $i \neq j$, and $\cup_i \mathbf{C}_i = \mathbf{S}$) by minimizing the following energy function:

$$E_{partition}(\{\mathbf{C}_i\}_{i=1}^n) = \sum_{i=1}^{n} E(\mathbf{C}_i), \tag{1}$$

where $E(\mathbf{C}_i)$ denotes the properness of clustering the underlying primitives of cluster i together, i.e. usually a shape proxy is associated and the fitting errors from the shape proxy to all underlying primitives are utilized.

3.1 Energy Based on Quadric Error Metrics (QEM)

QEM method [9] uses a point $\mathbf{v}_i$ as the shape proxy for cluster $\mathbf{C}_i$. The geometric primitives for a cluster are tangent planes, i.e. each point $\mathbf{p} \in \mathbf{C}_i$ is equipped with its normal direction $\mathbf{n}(\mathbf{p})$. The energy of cluster $\mathbf{C}_i$ is defined as the integral of the squared distance from the optimal proxy (point $\mathbf{v}_i$) to all these tangent planes in $\mathbf{C}_i$:

$$E_{qem}(\mathbf{C}_i) = \min_{\mathbf{v}_i} E_{qem}(\mathbf{C}_i, \mathbf{v}_i), \tag{2}$$

$$E_{qem}(\mathbf{C}_i, \mathbf{v}_i) = \iint_{\mathbf{p}\in\mathbf{C_i}} (\mathbf{n}(\mathbf{p})^\top (\mathbf{v}_i - \mathbf{p}))^2 \, d\mathbf{p}. \tag{3}$$

Note that a plane can be extended to infinity in 3D space, thus the point proxy can fit a non-local plane with zero penalty. The minimization of the QEM energy will guide the clustering results to disconnected (broken) patches since a point proxy will always "absorb" non-local planes that nearly pass through it (see Fig. 1). In this paper, we will use the term "broken patches" to describe the effect that the minimization of an energy function naturally leads the clustering to broken patches.

Due to its lacking penalty on the disconnectedness, QEM method [9] only allows local edge contraction operation to

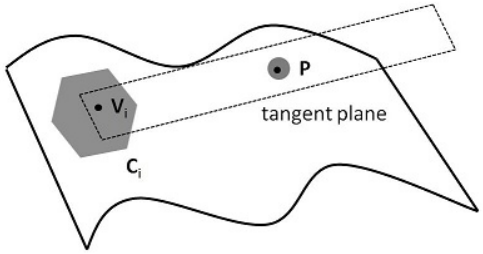

Fig. 1: Problems of clustering with QEM energy.

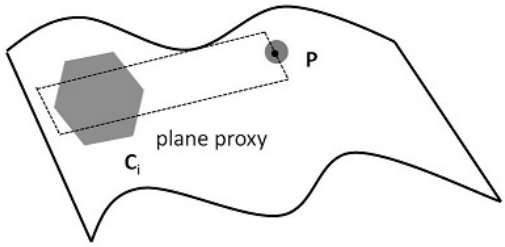

Fig. 2: Problems of clustering with plane fitting energy.

ensure each cluster to be connected. Each edge contraction aggregates the two corresponding clusters into a connected one and computes the optimal point for it. Even though the asymptotically optimal anisotropic behavior of triangle meshes generated by QEM has been proved [20], the greedy nature of edge contraction inevitably leads to suboptimal geometric partition result for non-infinitesimal clusters, as illustrated in Fig. 3.

3.2 Energy Based on Plane Fitting

The plane fitting method is widely accepted in computer graphics community through the work of Dual QEM (DQEM) [21] and VSA [8]. Note that besides the plane fitting energy, VSA [8] also proposed the $L^{2,1}$ shape metric, i.e. the L^2 measure between the normal field and the plane's unit normal. Even though the $L^{2,1}$ shape metric is recommended to be used by VSA [8], we focus our discussion on its plane fitting energy, since the $L^{2,1}$ shape metric does not lead the regions to follow the asymptotic anisotropic requirement of approximation theory and thus is not theoretically "on the right track" to achieve the optimal surface approximation, as can be seen from the comparison results of Fig. 12 and Fig. 13 later.

The plane fitting method uses a plane $\mathbf{L}_i$ as the shape proxy for each cluster $\mathbf{C}_i$. We denote the plane as $\mathbf{L}_i = (\mathbf{n}_i, d_i)^\top$ with $\mathbf{n}_i$ and d_i being its unit normal and scalar offset correspondingly. The energy of cluster $\mathbf{C}_i$ is defined as the integral of the squared distances from all points $\mathbf{p} \in \mathbf{C}_i$ to the best fitting plane:

$$E_{plane}(\mathbf{C}_i) = \min_{\mathbf{L}_i} E_p(\mathbf{C}_i, \mathbf{L}_i), \tag{4}$$

$$E_{plane}(\mathbf{C}_i, \mathbf{L}_i) = \iint_{\mathbf{p}\in\mathbf{C_i}} (\mathbf{n}_i^\top \mathbf{p} + d_i)^2 \, d\mathbf{p}. \tag{5}$$

This plane fitting approach suffers from the "broken patches" problem as the QEM energy : a plane proxy will always "absorb" non-local points it passes through (see Fig. 2). This problem is noticed in Sec. 3.5 of VSA paper [8]: convergence is not guaranteed if the connectedness of the partition regions are required.

To practically apply this energy for geometry partition, Garland et al. [21] used the greedy edge collapse as in QEM [9] to seek the suboptimal partition result. VSA [8] guarantees the connectedness of each cluster by performing the partition process triangle-by-triangle using a region flooding procedure driven by a priority queue, which contains new growing triangle candidates sorted by the error to their corresponding fixed plane proxy. Their final result is the nearly settled down partition of Lloyd's relaxation [13] on the two phases of partitioning and plane fitting. It is worth noting that even though the flooding scheme guarantees the connectedness, however, a cluster may totally reside in another cluster for the plane fitting energy, as illustrated in Fig. 3 (It does not happen for the $L^{2,1}$ shape metric). As one would expect, this plane fitting energy prefers to partition the input model with a set of parallel planes along its least length direction.

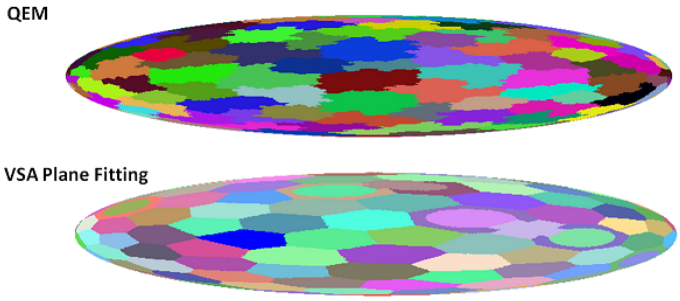

Fig. 3: Partition the Ellipsoid surface into 200 clusters using the QEM method [9] and the VSA plane fitting method [8]. Note that a cluster may totally reside in another one for the VSA plane fitting.

3.3 Energy Based on Other Proxies

Besides the point proxy and the plane proxy, other non-planar proxies such as spheres, cylinders and rolling ball blend patches [22], and quadric proxies [23] have been used for robust surface structure extraction. Even though these methods fit a quadric surface to a given cluster $\mathbf{C}_i$ in L^2 sense, there is no evidence that the anisotropy of their clusters has any relation to the asymptotic optimal anisotropy. Also these methods cannot be easily extended for remeshing purpose, thus discussing these methods is beyond the scope of this paper.

4 Our Energy Definition

In this section, we introduce a new error-driven energy whose minimization not only suppresses the clusters from being broken, but also naturally leads clusters to have the optimal aspect ratio and cluster size as specified in approximation theory [2] [3] [4] for L^∞ piecewise linear approximation to the given surface.

To better illustrate the properties of our proposed energy, in Sec. 4.1, we first introduce an overall scalar measurement of the cluster's variance, which is the covariance matrix's determinant. Then we show that this overall scalar variance has a CVT nature of suppressing the clusters from being broken. In Sec. 4.3, we formally formulate our energy and show how it relates to L^∞ piecewise linear approximation to

the given surface [2] [3] [4]. Last but not least, the degenerate surfaces (i.e. either one or both of the principal curvatures are degenerated to zero) are handled in Sec. 4.4.

4.1 Ellipsoidal Variance Model via PCA

When the plane is used as the shape proxy, the plane normal and scalar offset can be optimized effectively using PCA: the best fitting plane will pass through the cluster's centroid $\mathbf{x}_c(\mathbf{C}_i)$; its normal direction will be the eigenvector of the smallest eigenvalue of the covariance matrix $\mathbf{U}(\mathbf{C}_i)$. The details of their derivation are given in the supplementary material. Here the centroid $\mathbf{x}_c(\mathbf{C}_i)$ and the covariance matrix $\mathbf{U}(\mathbf{C}_i)$ are given by:

$$\mathbf{x}_c(\mathbf{C}_i) = \frac{\iint\limits_{\mathbf{p}\in\mathbf{C}_i} \mathbf{p}\, d\mathbf{p}}{\iint\limits_{\mathbf{p}\in\mathbf{C}_i} d\mathbf{p}}, \tag{6}$$

$$\mathbf{U}(\mathbf{C_i}) = \iint\limits_{\mathbf{p}\in\mathbf{C}_i} (\mathbf{p}-\mathbf{x}_c(\mathbf{C}_i))(\mathbf{p}-\mathbf{x}_c(\mathbf{C}_i))^\top\, d\mathbf{p}. \tag{7}$$

Based on the covariance matrix $\mathbf{U}(\mathbf{C}_i)$, an ellipsoidal variance model can be used to describe the variability of the cluster $\mathbf{C}_i$:

$$(\mathbf{x}-\mathbf{x}_c(\mathbf{C}_i))^\top \mathbf{U}(\mathbf{C}_i)^{-1}(\mathbf{x}-\mathbf{x}_c(\mathbf{C}_i)) = 1. \tag{8}$$

This ellipsoid has three axes $\mathbf{e}_j, j = 1,2,3$, which are the three eigenvectors of $\mathbf{U}(\mathbf{C}_i)$. The corresponding lengths of semi-principal axes are $l_j, j = 1,2,3$, which are the square roots of the corresponding eigenvalues ($\lambda_j, j = 1,2,3$): $l_j = \sqrt{\lambda_j}$. Note the length of l_j is the variance of the points in $\mathbf{C}_i$ measured in the corresponding $\mathbf{e}_j$ direction.

The plane fitting energy only considers the variance along the shortest semi-principal axis of the ellipsoidal variance model. The cause of the "broken patches" problem is due to the following fact: to minimize the plane fitting error, a cluster prefers to squeeze along the shortest semi-principal axis direction at the expense of unreasonably stretching along other directions. Inspired by the idea that a scalar error taking the variance in all directions into consideration may be an effective way to suppress the "broken patches" problem, we use the determinant of the covariance matrix. Note this idea has also been introduced as *generalized variance* for the scalar measurement of overall multidimensional data [24].

The covariance matrix's determinant can be interpreted as the squared volume of the ellipsoidal variance model multiplied by a constant factor:

$$\begin{aligned} det(\mathbf{U}(\mathbf{C}_i)) &= \prod_{j=1}^{3} \lambda_j \\ &= \frac{9}{16\pi^2}(\frac{4\pi}{3}\prod_{j=1}^{3}\sqrt{\lambda_j})^2 \\ &= \frac{9}{16\pi^2} \times \text{squared ellipsoid volume}. \end{aligned} \tag{9}$$

The "broken patche" problem will be suppressed using the squared volume of the ellipsoidal variance model. Classifying a point $\mathbf{p}$ far away from the centroid into $\mathbf{C}_i$ will result in a large increase for the variance along the direction of $\mathbf{p}-\mathbf{x}_c(\mathbf{C}_i)$, while the variances in other orthogonal directions keep unchanged. Thus it will cause a large increase for this energy. The minimization of this energy prefers the clustering to be one locally-connected patch. Actually, this measurement has a CVT nature of suppressing the clusters from being broken (See Sec.4.2).

4.2 Relation to Mahalanobis CVT

In this section, we show the connection between the determinant of the covariance matrix with the Mahalanobis CVT (MCVT) [25].

For the energy of each individual cluster, Richter et al. [25] used the integral of the Mahalanobis distance from the centroid to the points w.r.t. a metric to be optimized. To avoid the trivial zero matrix while optimizing the metric tensor, the metric is constrained to have a unit determinant. This constraint directly leads to the optimal metric being the inverse covariance matrix scaled by a factor, which is the d-th root (d is the space dimension) of the determinant of the covariance matrix (See Eq.(15) of MCVT paper [25]). Thus the cluster energy of MCVT is essentially the d-th root of the determinant of covariance matrix. The covariance matrix's determinant has the CVT nature to suppress broken clusters.

Even though MCVT models the cluster energy using a similar expression form with the covariance matrix's determinant, their way to evaluate the cluster energy is fundamentally different: MCVT utilizes the concept of optimized metric while we propose it from the perspective of variance analysis. More specifically, to evaluate the energy of a cluster, MCVT solves the optimized metric first, then it sums over the individual Mahalanobis distances over all candidate points. From our perspective, this inefficient computation can be replaced by a direct evaluation of the determinant from the covariance matrix.

4.3 Optimal Surface Approximation

In this section, we first review the desired asymptotic behavior for L^∞ surface approximation from the perspective of approximation theory. Then we show how these objectives are fulfilled by the minimization of the proposed PCA-based energy.

For the general L^p metric, the optimal piecewise linear approximation of a $\mathbf{C}^2$ surface will have the asymptotic stretching ratio as the ratio of the reciprocal square root of the two principal curvatures, which is $\sqrt{|k_{max}/k_{min}|}$, where k_{max} and k_{min} are the two principal curvatures [2], [3], [4]. For different p values in the L^p metric, the optimal surface approximations require different size allocation schemes on these linear elements to equidistribute the error. Here we focus our discussion on the L^∞ surface approximation since it is closely related to the Laplace-Beltrami operator and can be measured by the Hausdorff distance in practice. Appendix A shows that the optimal asymptotic cluster size is inversely proportional to the square root of the Gaussian curvature for error equidistribution.

Before introducing our energy definition, we illustrate an important property of the overall variance measure in Eq.(9): assume the given surface is $\mathbf{C}^2$, when the variance measure of the cluster is minimized for a fixed infinitesimal cluster size, the aspect ratio of the clusters will automatically converge to the optimal asymptotic stretching ratio $\sqrt{|k_{max}/k_{min}|}$ given by the hidden principal curvatures. This proof is shown in Appendix B.

Introducing a coefficient that is only related to the cluster area into the cluster energy formulation is an effective way to achieve the goal of implicit cluster size allocation without affecting the asymptotic anisotropic behavior. We propose the energy definition of each cluster to be its covariance matrix's determinant normalized by the fourth power of the area:

$$E_{PCA}(\mathbf{C}_i) = \frac{det(\mathbf{U}(\mathbf{C}_i))}{area^4(\mathbf{C}_i)}. \quad (10)$$

The minimization of the summation over all cluster energies naturally assigns the optimal relative size for each cluster in the asymptotic sense while preserving the desired anisotropy. The proof is illustrated in Appendix C. We denote this energy definition as the *PCA-based energy*.

In summary, given a $\mathbf{C}^2$ surface and a partition number n that is large enough, the minimization of our PCA-based energy over all clusters will lead to the partition result that each cluster has the stretching ratio $\sqrt{|k_{max}/k_{min}|}$ along the maximum principal curvature direction and has the relative size inversely proportional to the square root of the Gaussian curvature. This converged aspect ratio and cluster size are consistent with the optimal aspect ratio and cluster size of L^∞ piecewise linear interpolation approximation.

4.4 Planar Surface Handling

If the surface has pure planar regions, the PCA-based energy on these regions is always zero for any aspect ratio and region size, even if we set the partition as broken patches. In this case, seeking the minimum shape approximation error will not distinguish these failure results.

For planar surfaces, we set the goal to be obtaining equilateral triangles of the same size in these regions. Given a cluster, we first check whether the shape approximation error per unit area, which is $\frac{det(\mathbf{U}(\mathbf{C}_i))}{area^5(\mathbf{C}_i)}$, is smaller than a predefined small threshold t. If it is, then this cluster is a degenerate cluster. We extend our PCA-based energy using a quality measurement for the degenerate clusters:

$$E_{PCA}(\mathbf{C}_i) = \alpha \cdot trace(\mathbf{U}(\mathbf{C}_i)) \cdot area(\mathbf{C}_i), \quad (11)$$

where α is the quality measurement coefficient, which is a predefined small constant to balance the relative cluster size between the planar regions and non-planar regions. For a degenerate cluster, the minimization of the covariance matrix's trace leads to its first two eigenvalues being identical, thus resulting in the optimal aspect ratio being 1.

5 The Optimization Algorithm

In this section, we develop an effective discrete variational framework to seek the optimal partition for the PCA-based energy. In this discrete setting, clusters are built from discrete triangular faces and these triangular faces serve as the integral domain of the covariance matrix.

At first, each face of the given mesh is treated as an individual cluster. Our partition algorithm consists of two steps: the merging step and the swapping step. The merging step gives an initial partition of the surface into n clusters by successively merging neighboring cluster-pair till the user-assigned number of clusters n is achieved. The details are provided in Sec. 5.1. Greedily merging clusters will not optimize the partition energy since the merging operation restricts the faces in the two merged clusters to reside in the same final cluster. Thus we relax this binding through a swapping step. The optimal partition is achieved when no local face swapping could further decrease the energy. More details are discussed in Sec. 5.2. An illustrative example of our variational partition algorithm applied on the Kitten model is shown in Fig. 4.

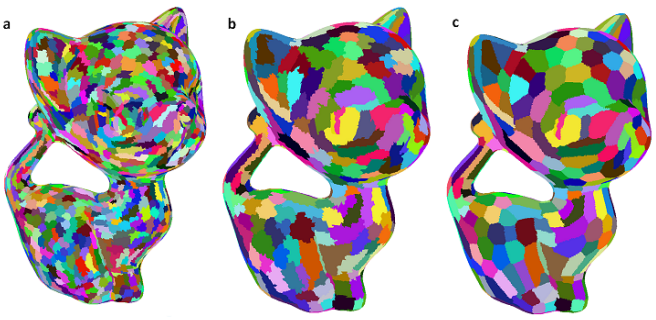

Fig. 4: Illustrative example for our variational merging-swapping algorithm: (a) Greedily apply merging operation to 2,000 clusters. (b) Successively merge till 500 clusters. (c) Optimal partition for 500 clusters after applying swapping operation.

Before introducing the main body of our algorithm, we would like to emphasize its efficiency, which makes the processing on complicated models up to tens of millions triangular faces practical. Our algorithm is very efficient due to two reasons. Firstly, the energy evaluation process for a cluster can be computed very efficiently. According to the energy definition of Eq. (10), we first evaluate the determinant of the covariance matrix. Then if it is a degenerate cluster, the trace of the covariance matrix will be computed. Both the determinant and the trace of a 3 by 3 matrix can be evaluated very efficiently without the need to perform eigen-decomposition. Secondly, maintaining these covariance matrices are computationally efficient for the merging operations and swapping operations. The supplementary material shows the computation of the covariance matrix for one triangle. The efficient update of the covariance matrix is described in Appendix D for the merging operation, and in Appendix E for the swapping operation. As can be seen from these appendices, all the computation needed to maintain the validity of the covariance matrix only involves several float number multiplications based on the shift of the cluster's centroid.

5.1 Initial Partition using Merging Operation

We utilize a merging operation for the initial partition. At first, each face is treated as a cluster, which may form 3 cluster-pairs with its neighboring faces that share common edges. Each cluster-pair can merge into a new cluster with a merging cost. That is, for a merging operation $(\mathbf{C}_i, \mathbf{C}_j) \rightarrow \mathbf{C}_k$, the cost is defined as the increase of the total energy: $E(\mathbf{C}_k) - E(\mathbf{C}_i) - E(\mathbf{C}_j)$ based on the definition of Eq. (10) or Eq. (12). As shown in Appendix D, the complexity of computing such merging cost is $O(1)$, which is independent of the number of triangles in each cluster. All possible merging operations are inserted into a heap, with the merging cost as the key factor. The heap always stores all possible merging operations for the current partition.

Then, the total number of clusters are reduced to n by successively merging the least-cost cluster-pair, which

is extracted from the top of the heap. After the least-cost pair merging is performed, only a local update is needed to maintain the validity of the heap: the remaining pairs of the two merged clusters are deleted, and the potential pairs of new clusters are computed and inserted. For each least-cost pair merging operation, the total number of clusters is reduced by 1 and the energy is expected to have the slowest increase rate among all the possible merging operations. This step is successively performed until the user-assigned cluster number n is achieved.

5.2 Partition Optimization using Swapping Operation

For partition optimization, we let each triangle face freely decide which cluster to reside among the final n clusters. For a triangle $\mathbf{T}$ residing in cluster $\mathbf{C}_i$, we can test whether swapping it to another cluster $\mathbf{C}_j$, where $j \neq i$, can decrease the energy in Eq. (1). In practice, to avoid the testing of unnecessary swapping, these swappings are launched in an iterative scheme. For each iteration, triangle $\mathbf{T}$ is limited to be the boundary triangles. These triangles can only swap to their neighboring clusters. This energy test only involves two clusters. Suppose after swapping, $\mathbf{C}_i$ becomes $\mathbf{C}_{i'}$, and $\mathbf{C}_j$ becomes $\mathbf{C}_{j'}$. That is, $\mathbf{C}_i = \mathbf{C}_{i'} \cup \mathbf{T}$, and $\mathbf{C}_{j'} = \mathbf{C}_j \cup \mathbf{T}$. The change of energy is $E(\mathbf{C}_{i'})+E(\mathbf{C}_{j'})-E(\mathbf{C}_i)-E(\mathbf{C}_j)$. As shown in Appendix E, the complexity of computing such energy change is $O(1)$, which is independent of the number of triangles in each cluster. Those swaps that decrease Eq. (1) will be accepted. The optimal partition is achieved when no swapping could further decrease the energy.

5.3 Partition Results

From theoretical perspective, our PCA-based energy guarantees the optimal approximation of clusters as linear elements for arbitrary $\mathbf{C}^2$ surface (Sec. 4.3), meanwhile the possible planar degenerate case is handled using the same PCA framework by the meshing quality measurement (Sec. 4.4). However, the triangular mesh models are only $\mathbf{C}^0$ continuous and may have sharp feature edges. In this section, we show the partition results with a finite partition number on several triangular mesh models using the discrete variational framework. We noticed not only the partition result is consistent with our theoretical derivation (See Fig. 5), it also can capture sharp features without any pre-tagging (See Fig. 6).

The partition on the Ellipsoid model is illustrated in Fig. 5(a). Note the "broken patches" problem shown in Fig. 3 is effectively suppressed due to the CVT nature of our PCA-based energy. Adaptation of the aspect ratio can be seen from the two ends of the ellipsoid to the middle part. The *relative aspect ratio error* is used to measure the anisotropy deviation:

$$\delta r = \frac{r_m - r_t}{r_t}, \tag{12}$$

where r_m is the measured aspect ratio as the squared root of the ratio between the covariance matrix's first two eigenvalues, r_t is the theoretic aspect ratio $\sqrt{|k_{max}/k_{min}|}$ at the projected cluster centroid. We list the relative aspect ratio error histogram of the Ellipsoid partition in Fig. 5(e). For the quantitative evaluation of equidistributing the L^∞ distance, the results in Fig. 12 and Fig. 13 are shown after the meshing step is introduced in Sec. 6.

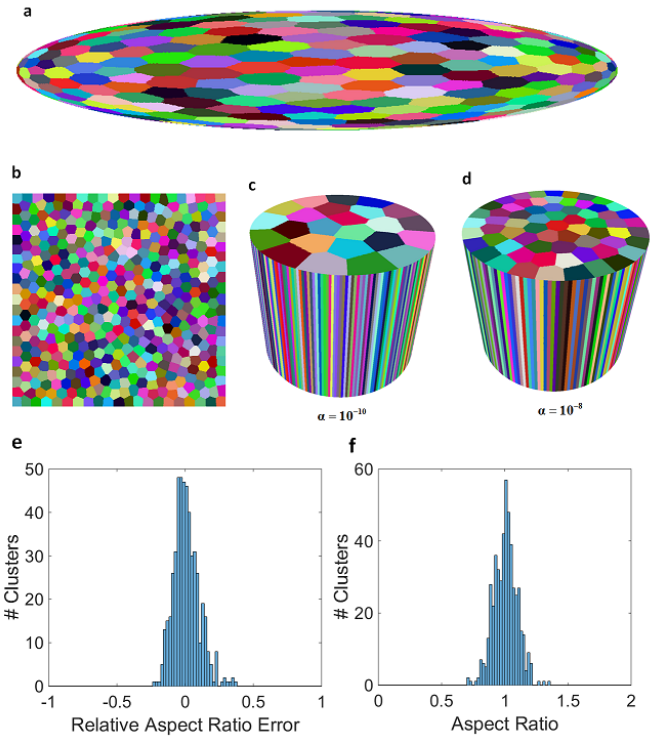

Fig. 5: Partition result on simple models: (a) Partition the Ellipsoid model into 500 clusters. (b) Partition the Plane model into 500 clusters. (c) and (d) Partition the Cylinder model into 300 clusters using different quality measurement coefficients α. (e) Relative aspect ratio error histogram of (a). (f) Aspect ratio histogram of (b).

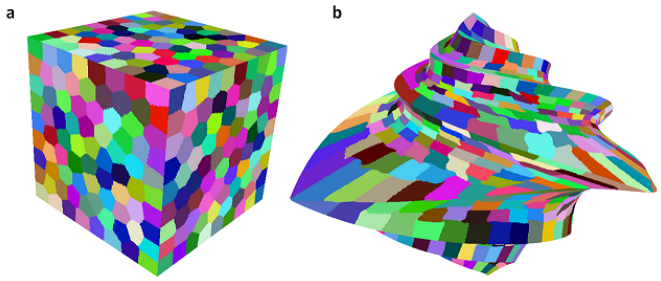

Fig. 6: Partition result on models with sharp feature: (a) Partition the Cube model into 500 clusters. (b) Partition the Octa-flower model into 1,000 clusters.

Fig. 5(b) shows the partition result on the Plane model with nearly hexagon tiling. Its aspect ratio histogram is listed in Fig. 5(f). Fig. 5(c) and Fig. 5(d) show the partition result on the Cylinder model. The Cylinder model has degenerate planar regions at its two ends while its side is non-degenerate since the shape approximation error is not zero for finite clusters. It can be seen that as the meshing quality measurement α increases in Fig. 5(d), the cluster size in the degenerate regions becomes smaller.

Fig. 6 clearly shows feature edges can be captured effectively by our PCA-based energy without any preprocessing. The shape approximation error is minimized by separating the clusters along these feature edges to avoid the sudden increment on the overall scalar variance measure. Fig. 8

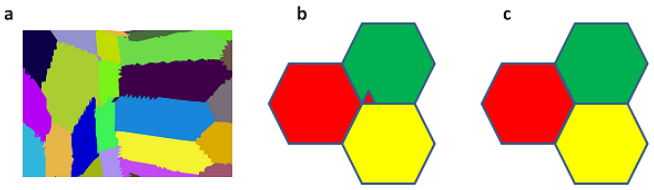

Fig. 7: The partition cleaning step: (a) The disconnectedness on real models. (b) and (c) illustrate merging the red cluster's disconnected component to its least cost neighbor

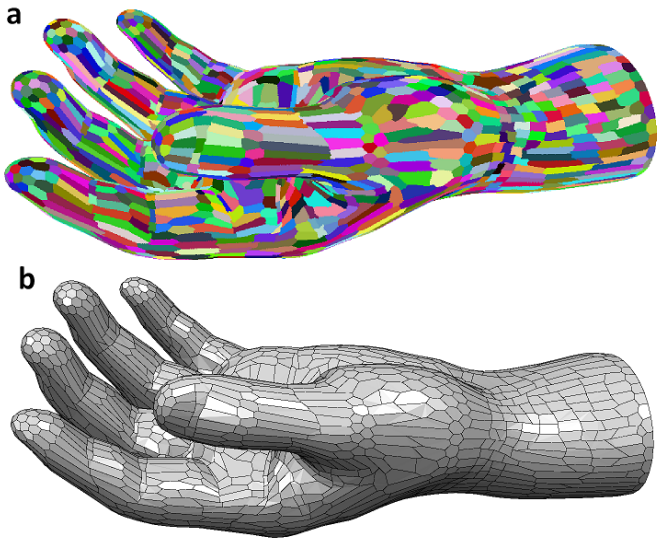

Fig. 8: Polygonal mesh generation on the Hand model: (a) Partition the Hand model into 2,000 clusters. (b) Extracted 2,000 polygons from the optimal partition.

illustrates the result on the Hand model.

6 Mesh Generation

As shown in Sec. 4.3, treating each cluster as a linear element maximizes the efficiency of approximation due to the asymptotic behavior of our optimized partition. It is straightforward to convert the optimal partition result into a polygonal mesh, the details are introduced in Sec. 6.1. If the linear elements are restricted to be triangles, we discuss the details of converting the polygonal mesh into a triangular one in Sec 6.2.

Before mesh generation, we perform a "partition cleaning" step [6] to guarantee each cluster to be a connected set of faces. We noticed disconnectedness do occur after our discrete optimization framework: since the basic testing unit of the swapping step is an entire triangle, the swapping of some triangles may be accepted with a very small energy decrease, while their neighbors are rejected due to a very small energy increase. In practice, usually dozens of boundary faces need to be cleaned on a model with several million faces. The disconnectedness is illustrated in Fig. 7(a). For these cases, the main component of the cluster is unchanged and all the "hanging components" can merge with their direct neighbors with the merging cost defined in Sec. 5.1. We merge these components to their least cost neighbor.

6.1 Polygonal Mesh Generation

Based on the optimal partition, our polygonal mesh generation algorithm is in the same spirit as VSA [8]: First, each cluster is attached with a plane proxy defined by the centroid and the smallest eigenvector direction of the corresponding cluster. These planes together serve as the optimal surface approximation. Secondly, anchor vertices are created to represent the original vertices where three or more regions meet. For each anchor vertex, its spatial position is set as the average of all the projections from the original vertices onto its neighboring plane proxies. Thirdly, the anchor vertices are connected based on the neighboring cluster information and are sorted into a polygon with respect to the original mesh's orientation.

The extracted polygonal meshes from the Ellipsoid surface and the Hand surface are shown in Fig. 9(a) and Fig. 8(b) respectively. Even though each polygon is not necessarily flat since the anchor vertex is the average of all projections. As reported in the VSA paper [8], this vertex enhancement technique generally improves the approximation accuracy of polygons and each polygon is indeed nearly flat. Also, it preserves the open boundaries on the given surface as shown in Fig. 8(b).

6.2 Triangular Mesh Generation

For triangular mesh generation, previous works [6] [10] usually treat the clustering result as the dual graph for triangular mesh. To generate a triangular mesh, one optimal vertex is assigned for each cluster and the connectivity of constructed mesh is defined by the adjacency information of the clustering. However, utilizing a dualization step directly is not applicable for our partition result. Our partition cuts the input models along its feature edges (See Fig. 6) and this dualization step will inevitably lose the control of the approximation of feature regions.

Our triangular mesh is built upon the polygonal mesh generated in Sec. 6.1. The polygonal mesh can be coarsely regarded as the optimal linear approximation of the given surface for a finite partition number, but they fail to satisfy the constraint of being triangular linear elements. We employ the discrete Constrained Delaunay Triangulation (CDT) step [8] to triangulate the polygons. Delaunay-like triangles are created while constraining the existing edges between anchor vertices to be in the output mesh. Then we remove some less geometrically important details using the technique of QEM [9] to reach the desired number of vertices or faces.

Fig. 9 shows the triangular mesh generation process on the Ellipsoid model with 500 assigned vertex budget. Fig. 10 illustrates the triangular mesh generation process on the Octa-flower model with $1,000$ assigned triangle budget. For both cases, we optimize the partition with the user assigned number and terminate the QEM contraction on the CDT mesh when this assigned condition is satisfied.

Fig. 11 shows two more complex examples on the Armadillo model and the Dragon model. Note the input mesh of the Armadillo model with a dense sampling rate with around 7.7M vertices and 15.5M faces is simplified into $5,000$ faces with nicely preserved details. The computational time is listed in Tab. 2.

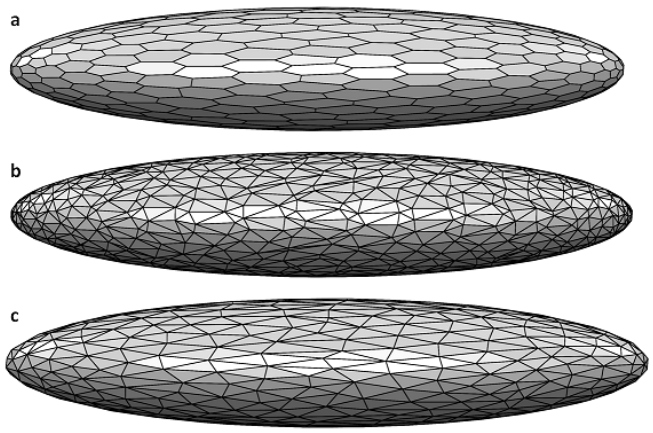

Fig. 9: Triangular mesh generation on the Ellipsoid model: (a) 500 polygons. (b) The CDT mesh with 1,988 triangles. (c) The final mesh with 500 vertices after QEM simplification.

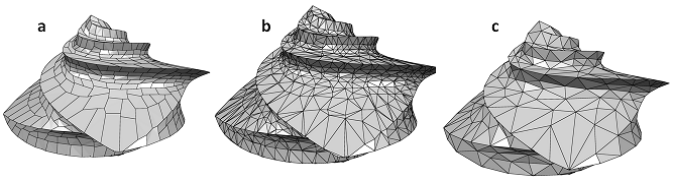

Fig. 10: Triangular mesh generation on the Octa-flower model: (a) 1,000 polygons. (b) The CDT mesh with 3,945 triangles. (c) The final mesh with 1,000 faces after QEM simplification.

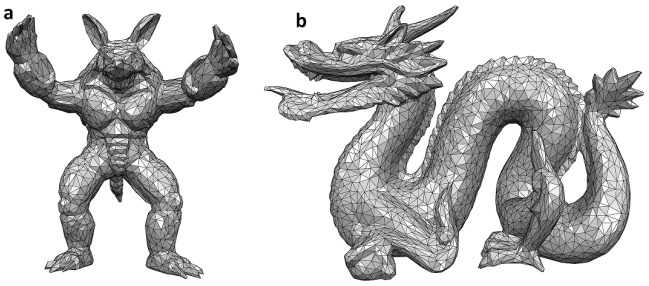

Fig. 11: Triangular mesh generation on other models: (a) 5,000 faces triangular mesh on the Armadillo model. (b) 10,000 faces triangular mesh on the Dragon model.

7 Experiments and Results

In this section, we demonstrate the effectiveness and efficiency of our method. Our algorithm is implemented using Microsoft Visual C++ 2010. For the hardware platform, the experiments are run on a desktop computer with Intel(R) Core i7-4770 CPU with 3.40GHz and 32GB DDR3 RAM.

To evaluate the quality for surface approximation, the *one-sided approximation error* from the original model to the remeshing results are computed using the Metro tool [26], i.e. for each vertex of the original mesh, the Metro tool searches its closest point on the generated triangular mesh. The *mean approximation error*, which is the mean value of these computed errors, is adopted as the criteria for the approximation quality. These error values are normalized w.r.t. the diagonal length of the original model's bounding box.

7.1 Polygonal Meshing for Function Approximation

For the error-driven approaches, the partition result and the corresponding proxies play the decisive role for the approximation ability. Using the polygonal mesh to represent the partition result and the proxies, we compare with the most widely used VSA approach [8]. Since the L^2 metric of VSA produce poor partition results (the same problem as shown in Fig. 3) after the plane proxies settle down from Lloyd's relaxation, the partition results of its recommended $L^{2,1}$ metric are used for comparison.

The input mesh is a dense sampling on a paraboloid function:

$$x^2 + y^2 - z = 0, -1 \leq x, y \leq 1,$$

with $263,169$ vertices and $524,288$ faces. Since the Hessian of this function is constant everywhere, the optimally approximation requires each partition region has the same region size with the aspect ratio of 1. Our partition result shows a nearly hexagonal tiling of the original surface, as shown in Fig. 12. We explicitly perform the discrete CDT step on both polygonal meshes to generated the flat surfaces for height difference evaluation. It can be seen that our method produces a smaller and more equidistributed height approximation error compared with VSA.

7.2 Polygonal Meshing for Surface Approximation

Fig. 13 shows the comparison on the Masque surface, which has $1,467,363$ vertices and $2,931,573$ faces. Here the approximation distance is evaluated at every vertices of the original mesh. We compute the one-sided approximation error to the polygonal mesh after the discrete CDT step. It clearly shows our method outperforms the VSA method with the one-sided approximation error nearly equidistribute everywhere. Tab. 1 lists the computational time of applying our approach for polygonal mesh generation on different models.

TABLE 1: Statistics for the polygonal mesh generation.

Model	# input vertex	# input face	# output polygon	merging time(s)	swapping time(s)	meshing time(s)	total time(s)
Ellipsoid	163,842	327,680	500	2.87	7.14	1.03	11.04
Paraboloid	263,169	524,288	500	5.29	21.01	1.59	27.89
Dragon	433,197	865,914	10,000	9.01	9.17	4.01	22.19
Octa-flower	497,789	995,560	1,000	10.51	47.08	3.03	60.62
Buddha	1,224,513	2,449,022	20,000	28.95	33.06	10.95	72.96
Masque	1,467,363	2,931,573	1,000	41.63	139.59	9.70	190.93
Hand	2,201,908	4,402,830	5,000	71.95	178.12	17.55	267.62

7.3 Triangular Meshing for Surface Approximation

In order to evaluate the effectiveness of our triangular mesh, we compare our results with Valette et al's discrete ACVT approach [6], Fu et al's LCT method [7] and the open-source remeshing tool provided by MeshLab v1.0.2 [27], which combines the QEM method with some post-simplification cleaning and refinement. Note that the ACVT approach and the LCT approach require the curvature metric as input. ACVT uses Cazals and Pouget's approach [28] for curvature estimation and LCT adopts Rusinkiewicz's method [29] and applies the same techniques as Boissonnat et al. [5] to smooth the estimated curvature tensor.

The comparison is performed on the Kitten model and the Buddha model. To test the ability of different methods, we compare them on a set of vertex budgets. Fig. 14 shows the mean approximation error curve w.r.t. the number of

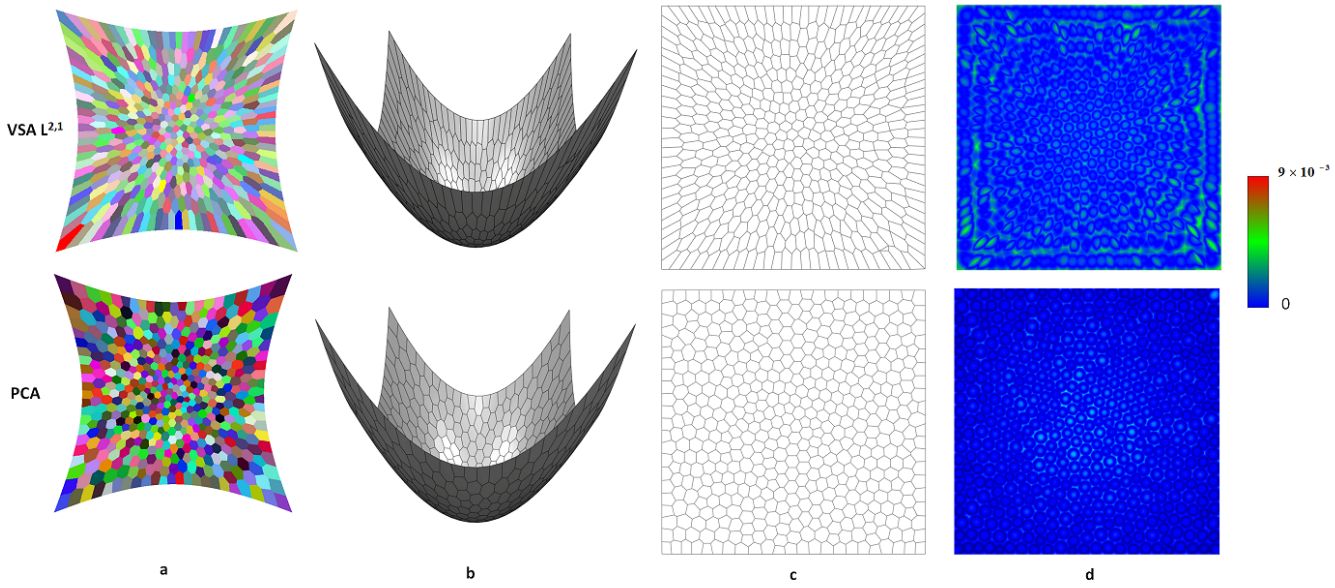

Fig. 12: Comparison of the polygonal meshes (500 polygons) on the paraboloid surface: (a) Partition result (top view). (b) The polygonal mesh (front view). (c) The polygons mapped onto the domain. (d) The height difference between the interpolated height of the CDT triangulation and the ground truth.

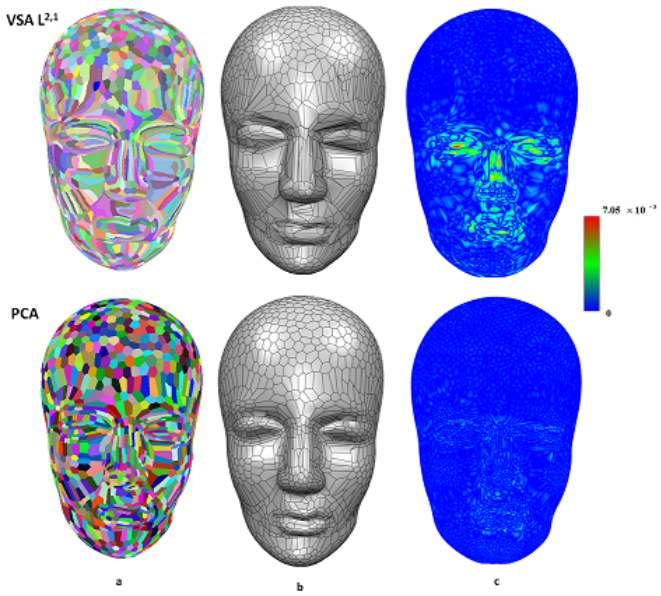

Fig. 13: Comparison of the polygonal meshes (1000 polygons) on the Masque model: (a) Partition result. (b) The polygonal mesh. (c) Color-encoded one-sided approximation error distribution on the original mesh.

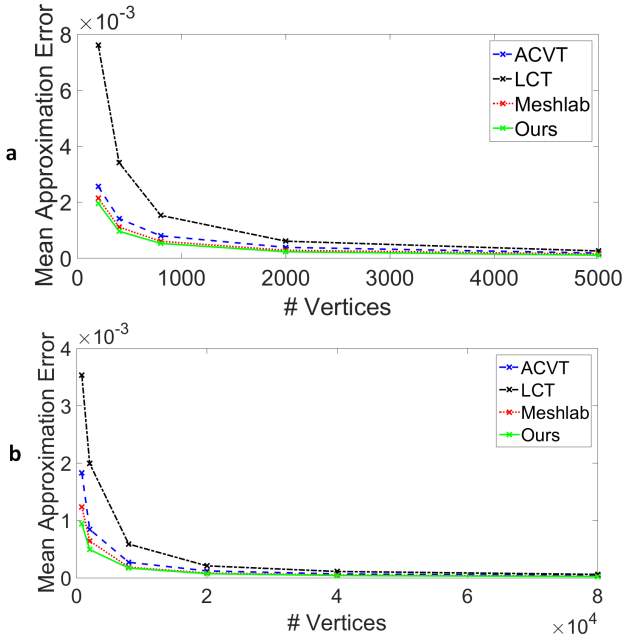

Fig. 14: The mean approximation error w.r.t. the number of vertices of remeshed models. (a) The Kitten model. (b) The Buddha model.

vertices. It can be clearly seen that our approach provides the best surface approximation among all these methods. Meshlab's approximation error is closest to our method. For the same approximation error, our method can save about $20\% \sim 30\%$ vertices compared with Meshlab. It should be noted the features of the models are not pre-tagged for the comparison, thus LCT method may miss the features during the projection step from the optimized vertices onto the triangular mesh, resulting in a much higher approximation error.

Fig. 15 shows the approximation result on the Kitten model with 200 vertex budget. Fig. 16 illustrates the approximation result on the Buddha model with $8,000$ vertex budget. Among these results, our method preserves the features (the neck and tail of the Kitten model, the pattern on the base and the fold on the skirt of the Buddha model), meanwhile exhibiting naturally adapted aspect ratio. For the LCT approach, obvious feature blurring can be noticed since the features are not identified in advance. The ACVT approach and the Meshlab tool show inferior feature align-

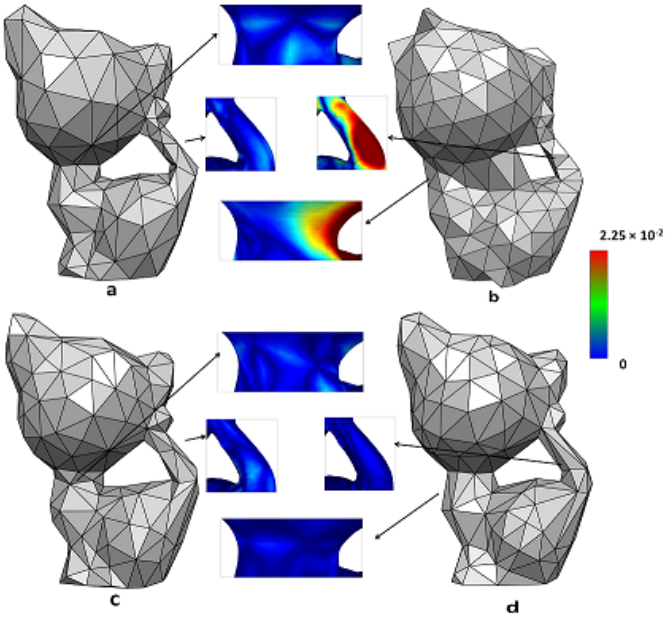

Fig. 15: Comparison on the Kitten model with 200 vertex budget: (a) The ACVT approach. (b) The LCT approach. (c) The Meshlab result. (d) Our result. Note the color-encoded error distribution is shown on the original model.

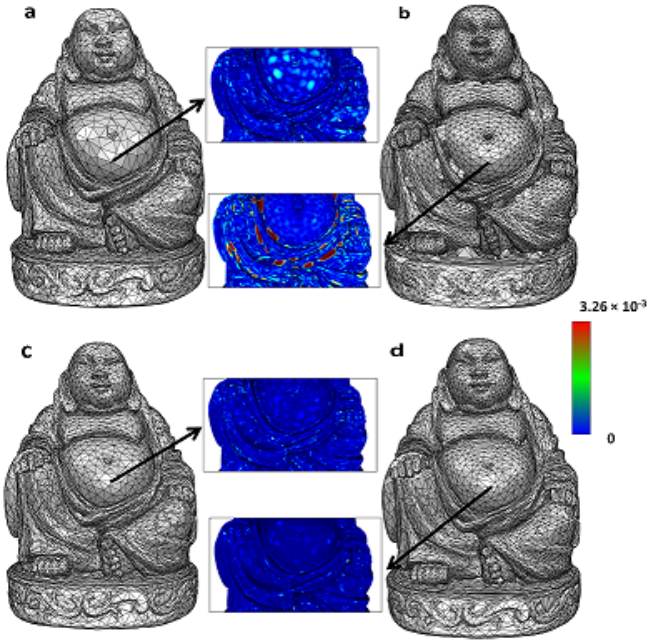

Fig. 16: Comparison on the Buddha model with 8,000 vertex budget: (a) The ACVT approach. (b) The LCT approach. (c) The Meshlab result. (d) Our result. Note the color-encoded error distribution is shown on the original model.

ment result compared with ours (See the neck of the Kitten model). In general, they put much denser vertices at the feature regions and coarser vertices at other regions. These results clearly show the advantageous of our PCA approach, i.e. being curvature-metric-free and feature awareness.

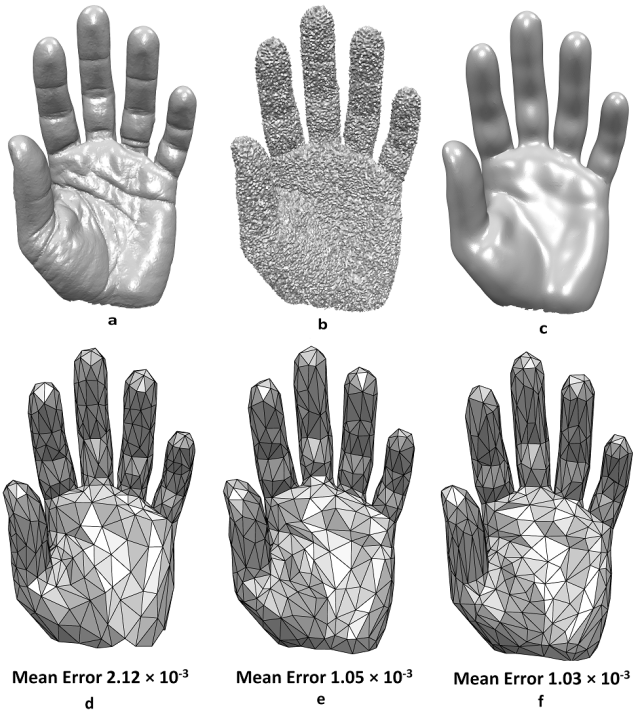

Fig. 17: The Olivier hand surface: (a) The original model. (b) The noisy model. (c) The denoised model after applying Sun et al.'s method [30]. (d), (e) and (f) 500 vertices remeshing results on the denoised model using ACVT, Meshlab, PCA respectively. The error is measured on the vertices of the original surface shown in (a).

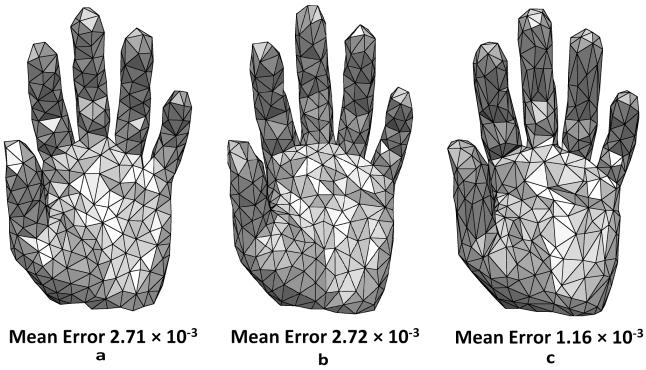

Fig. 18: The remeshing result on the noisy Olivier hand model using 500 vertices: (a) ACVT approach. (b) Meshlab result. (c) Our result. The error is measured on the vertices of the original model, shown in Fig. 17(a).

7.4 Triangular Meshing on Noisy Surface

Robustness against noises is usually desired for remeshing methods since the acquisition process of 3D scanners inevitably introduces noise signals. We model this acquisition process with the most widely used additive white Gaussian noise as illustrated in Fig. 17. Here the original Olivier hand model has $53,054$ vertices and $105,860$ faces and the noisy model is generated by disturbing the vertices of the original surface along their normal directions with Gaussian white noises whose variance is 0.002 w.r.t. the diagonal length of the original model's bounding box. With such severe noise mixed in, it is straightforward to apply some denoising

TABLE 2: Statistics for the triangular mesh generation.

Model	# input vertex	# input face	# output vertex	merging time(s)	swapping time(s)	meshing time(s)	total time(s)
Noisy Olivier hand	53,054	105,860	500	0.90	1.25	0.56	2.71
Kitten	137,098	274,196	200	2.74	9.72	1.21	13.67
Dragon	433,197	865,914	10,000	9.01	9.17	5.27	23.45
Octa-flower	497,789	995,560	1,000	10.51	47.08	4.26	61.85
Buddha	1,224,513	2,449,022	20,000	28.95	33.06	15.69	77.70
Armadillo	7,737,050	15,474,096	5,000	300.33	697.53	93.13	1070.82

techniques as preprocessing. As illustrated in Fig. 17(c), the denoising technique proposed by Sun et al. [30] removes the noise at the expense of filtering out important detail of the original model. Fig. 17(d), (e) and (f) show the remeshing results on the denoised model using ACVT, Meshlab and PCA respectively.

Note the grain details of the hand are lost during the denoising step and are unable to be recovered during the remeshing step. Thus we also discuss the alternative approach of remeshing on the challenging noisy surface directly. Fig. 18 shows the remeshing results on the noisy model shown in Fig. 17(b) with 500 vertices using ACVT, Meshlab and PCA. Among these results, our method provides the best visual result and still exhibits the anisotropy under severe additive noise. ACVT method requires the explicitly extraction of the curvature field, which quickly degenerates to meaningless random field under severe noise. Thus their anisotropic characteristic for better surface approximation is not robust against noise. The QEM result shown in Fig. 18(b) basically suffers from the same problem: the normal of triangle plane is not robust enough under noise attack thus their energy is also fragile. While our PCA-based approach is more stable to capture the anisotropy.

8 Conclusion & Discussion

As one of the error-driven approaches, our PCA-based energy is closely linked with the CVT energy thus penalizes the "broken patches". Besides the efficiency for energy evaluation, its anisotropic property and automatically capturing sharp features are well suited for surface remeshing.

Using a discrete variational optimization framework, the number of triangles of input mesh and its quality will influence our remeshing result. Generally speaking, the output from 3D scanners can be used directly. For triangular mesh with insufficient samples, subdivision of triangles is needed as preprocessing to obtain a denser input mesh.

Implicitly incorporating PCA of surface patches is a powerful way to provide a concise geometric representation, which is well demonstrated through the examples in our paper. Moreover, it seems to offer interesting alternatives for other types of approximation problems, such as the surface approximation problem arising from point cloud data, the height field approximation problem coming from images, etc. We will investigate these problems in the future.

Appendix A
Desired Asymptotic Cluster Size for L^∞ Approximation

For an infinite partition number n that each cluster converges to a patch with zero area on a $\mathbf{C}^2$ surface, we now derive the relationship between the cluster size and Gaussian curvature such that the L^∞ approximation error is equidistributed everywhere.

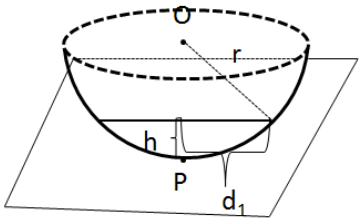

Fig. 19: The osculating circle along one principal direction.

We first consider the desired cluster length along one principal direction (see Fig. 19). For point $\mathbf{p}$ on the surface, its osculating circle, whose center is denoted as $\mathbf{O}$, has the radius along this direction as the inverse of the corresponding curvature $r = \frac{1}{k_1}$. We denote the equidistributed L^∞ approximation error as h. So for the infinitesimal patch of $\mathbf{p}$, h is the distance from point $\mathbf{p}$ to the linear interpolation element (it is the largest distance in this cluster), whose length at this specific direction is denoted as $2d_1$. It is obvious to see from the figure that h and d_1 must satisfy:

$$h = r - \sqrt{r^2 - d_1^2} = \frac{d_1^2}{r + \sqrt{r^2 - d_1^2}}. \tag{13}$$

When the number of clusters grows to infinity, we have $d_1 \to 0$. Thus $h \to \frac{d_1^2}{2r}$. Since h is identically distributed everywhere, we have $d_1 \sim \sqrt{r} \sim \frac{1}{\sqrt{k_1}}$.

For the other principal direction, this relationship still holds : $d_2 \sim \frac{1}{\sqrt{k_2}}$. So the optimal asymptotic cluster size that equidistributes the L^∞ approximation error is inversely proportional to the square root of Gaussian curvature: $\epsilon_{desire}^2 \sim d_1 d_2 \sim \frac{1}{\sqrt{k_1 k_2}}$.

Appendix B
Asymptotically Optimal Aspect Ratio

For an infinite partition number n that each cluster converges to a patch with zero area on a $\mathbf{C}^2$ surface, we now prove that the minimization of the overall variance measure, which is the determinant of the covariance matrix, will lead these infinitesimal patches to have the optimal aspect ratio $\sqrt{|\frac{k_2}{k_1}|}$.

Since the aspect ratio of a patch is involved, for the sake of simplicity, we assume that each patch is a rectangular shape whose sides are parallel to the principal curvatures, and have lengths ϵ_i, $i = 1, 2$ (see Fig. 20). The covariance matrix of this patch only depends on the two side lengths for the integration domain. We now investigate the optimal aspect ratio of $\frac{\epsilon_1}{\epsilon_2}$ to minimize the overall variance measure, which is determinant of the cluster's covariance matrix.

Let $\mathbf{C}$ be the infinitesimal rectangular patch on the surface, which has two principal curvatures k_i, $i = 1, 2$. Denote $\epsilon = \sqrt{\epsilon_1 \epsilon_2}$, and the aspect ratio of the rectangle as $r = \frac{\epsilon_1}{\epsilon_2}$. The covariance matrix $\mathbf{U}(\mathbf{C})$ is the surface integral over the rectangular domain:

$$\mathbf{U}(\mathbf{C}) = \mathbf{U}(\epsilon, r) = \iint_{\mathbf{p} \in \mathbf{C}} (\mathbf{p} - \mathbf{x}_c)(\mathbf{p} - \mathbf{x}_c)^\top \mathrm{d}\mathbf{p}, \tag{14}$$

where $\mathbf{p}$ is the points in the rectangle domain, $\mathbf{x}_c$ is the centroid of the rectangular patch:

$$\mathbf{x}_c = \frac{\iint_{\mathbf{p} \in \mathbf{C}} \mathbf{p} \mathrm{d}\mathbf{p}}{\iint_{\mathbf{p} \in \mathbf{C}} \mathrm{d}\mathbf{p}}. \tag{15}$$

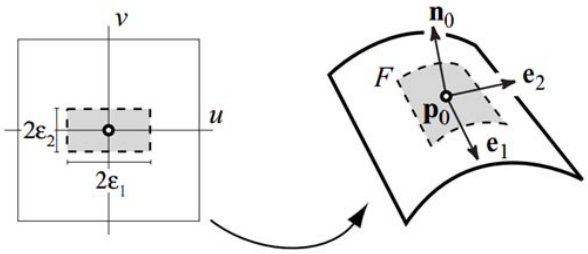

Fig. 20: Local parameterization of point $\mathbf{p}_0$ on the surface, $\mathbf{u}$, $\mathbf{v}$ axes coincide with the principal axes $\mathbf{e}_1$ and $\mathbf{e}_2$.

The overall covariance measure thus is $\prod_{j=1}^{3}\lambda_j(\mathbf{U}(\mathbf{C}))$. We are going to prove that when each rectangular patch converges to zero-area, the aspect ratio of each cluster will automatically converge to $\sqrt{\left|\frac{k_2}{k_1}\right|}$ to minimize its overall covariance measure:

$$r^\star = \operatorname*{argmin}_r \lim_{\epsilon\to 0}\prod_{j=1}^{3}\lambda_j(\mathbf{U}(\epsilon,r)) = \sqrt{\left|\frac{k_2}{k_1}\right|}. \tag{16}$$

Our proof first gives the formulas of the three eigenvalues of $\mathbf{U}(\epsilon,r)$ in terms of ϵ and r explicitly, then the best aspect ratio $r^\star$ that minimize the overall variance measure is derived based on these formulas.

The neighborhood of a point $\mathbf{p}_0$ on the surface can be approximated to the second order by:

$$\mathbf{p}(u,v) \approx [u, v, \frac{k_1u^2+k_2v^2}{2}]^\top, \tag{17}$$

where u and v are the local paramters as shown in Fig. 20.

We first consider the centroid $\mathbf{x}_c$:

$$\mathbf{x}_c = \frac{\iint_{\mathbf{p}\in\mathbf{C}}\mathbf{p}\mathrm{d}\mathbf{p}}{\iint_{\mathbf{p}\in\mathbf{C}}\mathrm{d}\mathbf{p}} = \frac{\iint_{\mathbf{p}\in\mathbf{C}}\begin{pmatrix} u \\ v \\ \frac{1}{2}(k_1u^2+k_2v^2)\end{pmatrix}\mathrm{d}\mathbf{p}}{\iint_{\mathbf{p}\in\mathbf{C}}\mathrm{d}\mathbf{p}} = \begin{pmatrix} 0 \\ 0 \\ p_{c3}\end{pmatrix}, \tag{18}$$

$$\begin{aligned}\iint_{\mathbf{p}\in\mathbf{C}}\mathrm{d}\mathbf{p} &\approx \iint_{\mathbf{p}\in\mathbf{C}}\sqrt{1+k_1^2u^2+k_2^2v^2}\mathrm{d}\mathbf{p}\\ &= \iint_{\mathbf{p}\in\mathbf{C}}(1+\mathbf{O}(u^2)+\mathbf{O}(v^2))\mathrm{d}\mathbf{p}\\ &\approx 4\epsilon_1\epsilon_2,\end{aligned} \tag{19}$$

$$\begin{aligned}p_{c3} &\approx \frac{\iint_{\mathbf{p}\in\mathbf{C}}\frac{1}{2}(k_1u^2+k_2v^2)\sqrt{1+k_1^2u^2+k_2^2v^2}\mathrm{d}\mathbf{p}}{4\epsilon_1\epsilon_2}\\ &= \frac{\iint_{\mathbf{p}\in\mathbf{C}}\frac{1}{2}(k_1u^2+k_2v^2)+\mathbf{O}(u^4+u^2v^2+v^4)\mathrm{d}\mathbf{p}}{4\epsilon_1\epsilon_2}\\ &\approx \frac{\frac{2}{3}\epsilon_1\epsilon_2(k_1\epsilon_1^2+k_2\epsilon_2^2)}{4\epsilon_1\epsilon_2}\\ &= \frac{1}{6}(k_1\epsilon_1^2+k_2\epsilon_2^2).\end{aligned} \tag{20}$$

Then the covariance matrix becomes:

$$\mathbf{U}(\mathbf{C}) = \iint_{\mathbf{p}\in\mathbf{C}}(\mathbf{p}-\mathbf{x}_c)(\mathbf{p}-\mathbf{x}_c)^\top\mathrm{d}\mathbf{p} \tag{21}$$

$$\begin{aligned}&= \iint_{\mathbf{p}\in\mathbf{C}}\begin{pmatrix} u \\ v \\ w\end{pmatrix}\begin{pmatrix} u & v & w\end{pmatrix}\\ &\qquad \sqrt{1+k_1^2u^2+k_2^2v^2}\mathrm{d}u\mathrm{d}v,\end{aligned} \tag{22}$$

where $w = \frac{1}{2}(k_1u^2+k_2v^2) - \frac{1}{6}(k_1\epsilon_1^2+k_2\epsilon_2^2)$. All the non-diagonal terms of the covariance matrix are 0, since they are odd functions of u or/and v, and the three diagonal terms U_{11}, U_{22}, and U_{33} correspond to the three eigenvalues:

$$\lambda_1 = U_{11} = \iint_{\mathbf{p}\in\mathbf{C}}u^2(1+\mathbf{O}(u^2)+\mathbf{O}(v^2))\mathrm{d}u\mathrm{d}v \approx \frac{4}{3}\epsilon_1^3\epsilon_2, \tag{23}$$

$$\lambda_2 = U_{22} = \iint_{\mathbf{p}\in\mathbf{C}}v^2(1+\mathbf{O}(u^2)+\mathbf{O}(v^2))\mathrm{d}u\mathrm{d}v \approx \frac{4}{3}\epsilon_1\epsilon_2^3, \tag{24}$$

$$\begin{aligned}\lambda_3 = U_{33} &= \iint_{\mathbf{p}\in\mathbf{C}}w^2(1+\mathbf{O}(u^2)+\mathbf{O}(v^2))\mathrm{d}u\mathrm{d}v\\ &= \frac{4}{45}(k_1^2\epsilon_1^5\epsilon_2+k_2^2\epsilon_1\epsilon_2^5)\\ &= \frac{4}{45}(k_1^2\epsilon^6r^2+k_2^2\epsilon^6r^{-2}).\end{aligned} \tag{25}$$

Thus,

$$E_{PCA}(\epsilon,r) = \prod_{j=1}^{3}\lambda_j = \frac{64}{405}\epsilon^{14}(k_1^2r^2+k_2^2r^{-2}). \tag{26}$$

The best aspect ratio $r^\star$ can be achieved at its stationary point:

$$\frac{\partial(E_{PCA}(\epsilon,r))}{\partial r}|_{r=r^\star} = 0 \tag{27}$$

$$\implies r^\star = \sqrt{\left|\frac{k_2}{k_1}\right|}. \tag{28}$$

Appendix C
Asymptotically Optimal Cluster Size for L^∞ Approximation

Following the derivations in Appendix B, we now extend the discussion to our proposed PCA energy, which is the summation over all clusters w.r.t. the determinant of its covariance matrix normalized by the fourth power of the area:

$$\begin{aligned}E_{PCA}(\{\mathbf{C}_i\}_{i=1}^n) &= \sum_{i=1}^{n}\frac{det(\mathbf{U}(\mathbf{C}_i))}{area^4(\mathbf{C}_i)}\\ &= \sum_{i=1}^{n}\frac{64}{405}(\epsilon^i)^6((k_1^i)^2(r^i)^2+(k_2^i)^2(r^i)^{-2}),\end{aligned} \tag{29}$$

where the superscript i refer to the cluster index.

For each cluster, its energy is a function depending on its cluster size ϵ^i and its aspect ratio r^i. Since the surface is given and total cluster size is fixed, optimizing the partition energy becomes a constrained minimization problem. It is worth noting that the optimal aspect ratios r^i are independent of each other:

$$\frac{\partial(E_{partition}(\{\epsilon^i,r^i\}_{i=1}^n))}{\partial r^i} = \frac{\partial(E_{PCA}(\epsilon^i,r^i))}{\partial r^i}. \tag{30}$$

Thus the optimal aspect ratio of each cluster is the same as illustrated in Appendix B, the cluster energy of the i-th cluster becomes:

$$E_{PCA}(\epsilon^i) = \frac{128}{405}k_1^ik_2^i(\epsilon^i)^6. \tag{31}$$

For the minimization of our proposed partition energy when $n\to\infty$:

$$\begin{aligned}\text{min}\qquad &\sum_{i=1}^{n}k_1^ik_2^i(\epsilon^i)^6,\\ \text{subject to}\quad &\sum_{i=1}^{n}(\epsilon^i)^2 = A.\end{aligned}$$

where A is constant surface area.

We can use Lagrange multiplier to analyze the constrained minimization problem :

$$f(\{\epsilon^i\}_{i=1}^n) = \sum_{i=1}^n k_1^i k_2^i (\epsilon^i)^6 - \lambda(\sum_{i=1}^n (\epsilon^i)^2 - A). \quad (32)$$

For the optimal cluster size of i-th cluster, it must satisfy:

$$\frac{\partial f(\{\epsilon^i\}_{i=1}^n)}{\partial \epsilon^i}|_{\epsilon^i=\epsilon^{i\star}} = 6k_1^i k_2^i (\epsilon^{i\star})^5 - 2\lambda\epsilon^{i\star} = 0, \quad (33)$$

$$\frac{\partial f(\{\epsilon^i\}_{i=1}^n)}{\partial \lambda} = \sum_{i=1}^n (\epsilon^i)^2 - A = 0. \quad (34)$$

So λ is a constant and the optimal cluster size is inversely proportional to the square root of Gaussian curvature:

$$(\epsilon^{i\star})^2 = \sqrt{\frac{\lambda}{3k_1^i k_2^i}} \sim \frac{1}{\sqrt{k_1^i k_2^i}}. \quad (35)$$

This agrees with the desired asymptotic cluster size shown in Appendix A.

APPENDIX D
EFFICIENT UPDATE OF COVARIANCE MATRIX FOR MERGING OPERATION

For a merging operation $(\mathbf{C}_i, \mathbf{C}_j) \to \mathbf{C}_k$, we show that the update of the covariance matrix can be efficiently computed by keeping track of the centroid and surface area of the two clusters. That is, there is no need to sum over all of the vertices in these cluster.

For cluster $\mathbf{C}_i$, we denote its surface area as $A(\mathbf{C}_i)$ and its centroid as $\mathbf{x}_c(\mathbf{C}_i)$:

$$A(\mathbf{C}_i) = \iint_{\mathbf{p}\in\mathbf{C}_i} d\mathbf{p}, \quad (36)$$

$$\mathbf{x}_c(\mathbf{C}_i) = \frac{\iint_{\mathbf{p}\in\mathbf{C}_i} \mathbf{p}\, d\mathbf{p}}{A(\mathbf{C}_i)}. \quad (37)$$

Considering the merging operation $(\mathbf{C}_i, \mathbf{C}_j) \to \mathbf{C}_k$, it is obvious the surface area $A(\mathbf{C}_k)$ and the centroid $\mathbf{x}_c(\mathbf{C}_k)$ can be updated easily:

$$A(\mathbf{C}_k) = A(\mathbf{C}_i) + A(\mathbf{C}_j). \quad (38)$$

$$\mathbf{x}_c(\mathbf{C}_k) = \frac{A(\mathbf{C}_i)\mathbf{x}_c(\mathbf{C}_i) + A(\mathbf{C}_j)\mathbf{x}_c(\mathbf{C}_j)}{A(\mathbf{C}_i) + A(\mathbf{C}_j)}. \quad (39)$$

We consider the update of the covariance matrix $\mathbf{U}(\mathbf{C}_k)$, which is:

$$\mathbf{U}(\mathbf{C}_k) = \iint_{\mathbf{p}\in\mathbf{C}_k} (\mathbf{p} - \mathbf{x}_c(\mathbf{C}_k))(\mathbf{p} - \mathbf{x}_c(\mathbf{C}_k))^\top\, d\mathbf{p}. \quad (40)$$

The integral domain $\mathbf{C}_k$ can be split into two parts : $\mathbf{C}_k = \mathbf{C}_i \cup \mathbf{C}_j$. For the first part $\mathbf{p} \in \mathbf{C}_i$, substituting $\mathbf{p} - \mathbf{x}_c(\mathbf{C}_k)$ with $\mathbf{p} - \mathbf{x}_c(\mathbf{C}_i) + \mathbf{x}_c(\mathbf{C}_i) - \mathbf{x}_c(\mathbf{C}_k)$, it becomes:

$$\begin{aligned}&\iint_{\mathbf{p}\in\mathbf{C}_i} (\mathbf{p} - \mathbf{x}_c(\mathbf{C}_k))(\mathbf{p} - \mathbf{x}_c(\mathbf{C}_k))^\top\, d\mathbf{p}\\ &= \mathbf{U}(\mathbf{C}_i) + A(\mathbf{C}_i)(\mathbf{x}_c(\mathbf{C}_i) - \mathbf{x}_c(\mathbf{C}_k))(\mathbf{x}_c(\mathbf{C}_i) - \mathbf{x}_c(\mathbf{C}_k))^\top.\end{aligned} \quad (41)$$

The second part is alike, so we can update the covariance matrix $\mathbf{U}(\mathbf{C}_k)$ as:

$$\begin{aligned}\mathbf{U}(\mathbf{C}_k) &= \mathbf{U}(\mathbf{C}_i) + \mathbf{U}(\mathbf{C}_j)\\ &+ A(\mathbf{C}_i)(\mathbf{x}_c(\mathbf{C}_i) - \mathbf{x}_c(\mathbf{C}_k))(\mathbf{x}_c(\mathbf{C}_i) - \mathbf{x}_c(\mathbf{C}_k))^\top\\ &+ A(\mathbf{C}_j)(\mathbf{x}_c(\mathbf{C}_j) - \mathbf{x}_c(\mathbf{C}_k))(\mathbf{x}_c(\mathbf{C}_j) - \mathbf{x}_c(\mathbf{C}_k))^\top.\end{aligned} \quad (42)$$

APPENDIX E
EFFICIENT UPDATE OF COVARIANCE MATRIX FOR SWAPPING OPERATION

For a swapping operation which swaps a triangle face $\mathbf{T}$ from $\mathbf{C}_i$ to $\mathbf{C}_j$. Suppose after swapping, $\mathbf{C}_i$ becomes $\mathbf{C}_{i'}$ and $\mathbf{C}_j$ becomes $\mathbf{C}_{j'}$, i.e., $\mathbf{C}_i = \mathbf{C}_{i'} \cup \mathbf{T}$, and $\mathbf{C}_{j'} = \mathbf{C}_j \cup \mathbf{T}$.

It is obvious the surface areas and the centroids can be updated by:

$$A(\mathbf{C}_{i'}) = A(\mathbf{C}_i) - A(\mathbf{T}), \quad (43)$$

$$A(\mathbf{C}_{j'}) = A(\mathbf{C}_j) + A(\mathbf{T}), \quad (44)$$

$$\mathbf{x}_c(\mathbf{C}_{i'}) = \frac{A(\mathbf{C}_i)\mathbf{x}_c(\mathbf{C}_i) - A(\mathbf{T})\mathbf{x}_c(\mathbf{T})}{A(\mathbf{C}_i) - A(\mathbf{T})}, \quad (45)$$

$$\mathbf{x}_c(\mathbf{C}_{j'}) = \frac{A(\mathbf{C}_j)\mathbf{x}_c(\mathbf{C}_j) + A(\mathbf{T})\mathbf{x}_c(\mathbf{T})}{A(\mathbf{C}_i) + A(\mathbf{T})}. \quad (46)$$

This swapping operation can be interpreted as two merging operations: $\mathbf{C}_i$ is the cluster merged from $\mathbf{C}_{i'}$ and $\mathbf{T}$, and $\mathbf{C}_{j'}$ is the cluster merged from $\mathbf{C}_j$ and $\mathbf{T}$:

$$\begin{aligned}\mathbf{U}(\mathbf{C}_{i'}) &= \mathbf{U}(\mathbf{C}_i) - \mathbf{U}(\mathbf{T})\\ &- A(\mathbf{C}_{i'})(\mathbf{x}_c(\mathbf{C}_{i'}) - \mathbf{x}_c(\mathbf{C}_i))(\mathbf{x}_c(\mathbf{C}_{i'}) - \mathbf{x}_c(\mathbf{C}_i))^\top\\ &- A(\mathbf{T})(\mathbf{x}_c(\mathbf{T}) - \mathbf{x}_c(\mathbf{C}_i))(\mathbf{x}_c(\mathbf{T}) - \mathbf{x}_c(\mathbf{C}_i))^\top.\end{aligned} \quad (47)$$

$$\begin{aligned}\mathbf{U}(\mathbf{C}_{j'}) &= \mathbf{U}(\mathbf{C}_j) + \mathbf{U}(\mathbf{T})\\ &+ A(\mathbf{C}_j)(\mathbf{x}_c(\mathbf{C}_j) - \mathbf{x}_c(\mathbf{C}_{j'}))(\mathbf{x}_c(\mathbf{C}_j) - \mathbf{x}_c(\mathbf{C}_{j'}))^\top\\ &+ A(\mathbf{T})(\mathbf{x}_c(\mathbf{T}) - \mathbf{x}_c(\mathbf{C}_{j'}))(\mathbf{x}_c(\mathbf{T}) - \mathbf{x}_c(\mathbf{C}_{j'}))^\top.\end{aligned} \quad (48)$$

REFERENCES

[1] P. K. Agarwal and S. Suri, "Surface approximation and geometric partitions," *SIAM Journal on Computing*, vol. 27, no. 4, pp. 1016–1035, 1998.

[2] E. Nadler, "Piecewise linear best L_2 approximation on triangulations," in *Approximation Theory V*, C. Chui, L. Schumaker, and J. Ward, Eds. Academic Press, 1986, pp. 499–502.

[3] R. B. Simpson and R. B. Simpson, "Anisotropic mesh transformations and optimal error control," *Applied Numer. Math*, vol. 14, pp. 183–198, 1992.

[4] L. Chen, P. Sun, and J. Xu, "Optimal anisotropic simplicial meshes for minimizing interpolation errors in L^p-norm," *Mathematics of Computation*, vol. 76, no. 257, pp. 179–204, 2007.

[5] J.-D. Boissonnat, K.-L. Shi, J. Tournois, and M. Yvinec, "Anisotropic Delaunay meshes of surfaces," *ACM Trans. Graph.*, vol. 34, no. 2, pp. 14:1–14:11, 2014.

[6] S. Valette, J. M. Chassery, and R. Prost, "Generic remeshing of 3D triangular meshes with metric-dependent discrete Voronoi diagrams," *IEEE Transactions on Visualization and Computer Graphics*, vol. 14, no. 2, pp. 369–381, 2008.

[7] X.-M. Fu, Y. Liu, J. Snyder, and B. Guo, "Anisotropic simplicial meshing using local convex functions," *ACM Trans. Graph.*, vol. 33, no. 6, pp. 182:1–182:11, 2014.

[8] D. Cohen-Steiner, P. Alliez, and M. Desbrun, "Variational shape approximation," *ACM Trans. Graph.*, vol. 23, no. 3, pp. 905–914, 2004.

[9] M. Garland and P. S. Heckbert, "Surface simplification using quadric error metrics," in *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '97, 1997, pp. 209–216.

[10] Q. Du, V. Faber, and M. Gunzburger, "Centroidal Voronoi tessellations: Applications and algorithms," *SIAM Rev.*, vol. 41, no. 4, pp. 637–676, 1999.

[11] G. Leibon and D. Letscher, "Delaunay triangulations and Voronoi diagrams for Riemannian manifolds," in *Proceedings of the Sixteenth Annual Symposium on Computational Geometry*, ser. SCG '00, 2000, pp. 341–349.

[12] Q. Du and D. Wang, "Anisotropic centroidal Voronoi tessellations and their applications," *SIAM Journal on Scientific Computing*, vol. 26, no. 3, pp. 737–761, 2005.

[13] S. P. Lloyd, "Least squares quantization in PCM," *IEEE Transactions on Information Theory*, vol. 28, pp. 129–137, 1982.

[14] Y. Liu, W. Wang, B. Lévy, F. Sun, D.-M. Yan, L. Lu, and C. Yang, "On centroidal Voronoi tessellation – energy smoothness and fast computation," *ACM Trans. Graph.*, vol. 28, no. 4, pp. 101:1–101:17, 2009.

[15] B. Lévy and N. Bonneel, "Variational anisotropic surface meshing with Voronoi parallel linear enumeration," in *International Meshing Roundtable conf. proc.*, 2012, pp. 349–366.

[16] Z. Zhong, X. Guo, W. Wang, B. Lévy, F. Sun, Y. Liu, and W. Mao, "Particle-based anisotropic surface meshing," *ACM Trans. Graph.*, vol. 32, no. 4, pp. 99:1–99:14, 2013.

[17] L. Chen and J. Xu, "Optimal Delaunay triangulations," *Journal of Computational Mathematics*, vol. 22(2), pp. 299–308, 2004.

[18] J.-D. Boissonnat, C. Wormser, and M. Yvinec, "Locally uniform anisotropic meshing," in *Proceedings of the Twenty-fourth Annual Symposium on Computational Geometry*, ser. SCG '08, 2008, pp. 270–277.

[19] C. Dobrzynski and P. Frey, "Anisotropic Delaunay mesh adaptation for unsteady simulations," in *Proceedings of the 17th international Meshing Roundtable*, 2008, pp. 177–194.

[20] P. S. Heckbert and M. Garland, "Optimal triangulation and quadric-based surface simplification," *Journal of Computational Geometry: Theory and Applications*, vol. 14, pp. 49–65, 1999.

[21] M. Garland, A. Willmott, and P. S. Heckbert, "Hierarchical face clustering on polygonal surfaces," in *Proceedings of the 2001 Symposium on Interactive 3D Graphics*, ser. I3D '01, 2001, pp. 49–58.

[22] J. Wu and L. Kobbelt, "Structure recovery via hybrid variational surface approximation," *Computer Graphics Forum*, vol. 24, no. 3, pp. 277–284, 2005.

[23] D.-M. Yan, Y. Liu, and W. Wang, "Quadric surface extraction by variational shape approximation," in *Geometric Modeling and Processing - GMP 2006*, 2006, pp. 73–86.

[24] S. S. Wilks, *Mathematical statistics*. Wiley, 1962.

[25] R. Richter and M. Alexa, "Mahalanobis centroidal Voronoi tessellations," *Computers & Graphics*, vol. 46, no. 0, pp. 48–54, 2015.

[26] P. Cignoni, C. Rocchini, and R. Scopigno, "Metro: Measuring error on simplified surfaces." *Computer Graphics Forum*, vol. 17, no. 2, pp. 167–174, 1998.

[27] P. Cignoni, M. Callieri, M. Corsini, M. Dellepiane, F. Ganovelli, and G. Ranzuglia, "Meshlab: an open-source mesh processing tool." in *Eurographics Italian Chapter Conference*, V. Scarano, R. D. Chiara, and U. Erra, Eds. The Eurographics Association, 2008, pp. 129–136.

[28] F. Cazals and M. Pouget, "Estimating differential quantities using polynomial fitting of osculating jets." *Computer Aided Geometric Design*, vol. 22, no. 2, pp. 121–146, 2005.

[29] S. Rusinkiewicz, "Estimating curvatures and their derivatives on triangle meshes," in *Symposium on 3D Data Processing, Visualization, and Transmission*, 2004.

[30] X. Sun, P. Rosin, R. Martin, and F. Langbein, "Fast and effective feature-preserving mesh denoising," *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 5, pp. 925–938, 2007.